

SAP Business One, version for SAP HANA: Overview



Welcome to the overview of SAP Business One version for SAP HANA.

Objectives



At the end of this course, you will be able to:

- List the benefits of running SAP Business One, version for SAP HANA
- Explain the concepts behind the SAP HANA database

At the end of this course, you will be able to list the benefits of running SAP Business One on SAP HANA and explain some of the concepts underlying the SAP HANA database.

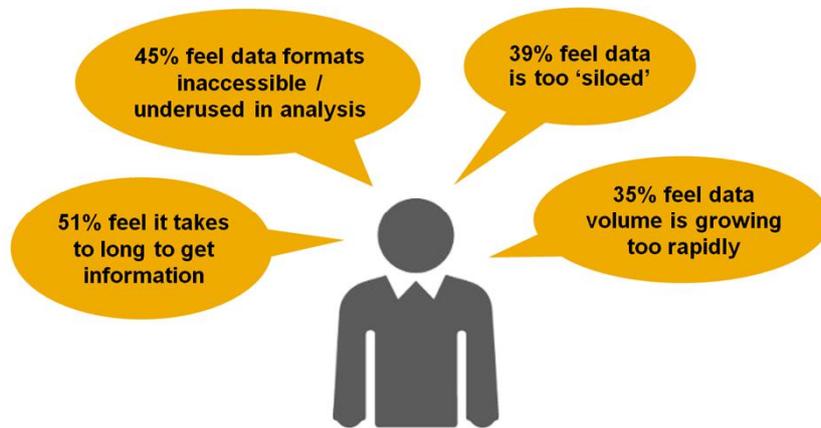


SAP HANA Overview



Let us discuss some basics about SAP HANA by looking at the business drivers which led to its development. Then we will take a high level look at the technology behind SAP HANA that enables us to transcend these issues. We bring it all together when we discuss the architecture underlying SAP Business One, version for SAP HANA.

SMEs feel the impact of big data



The "Big" aspect of Big Data is often in the eye of the beholder, and for smaller companies terabyte or even gigabyte scale data can pose the same problems and opportunities that petabytes do for global enterprises.

Aberdeen Group

© 2014 SAP AG. All rights reserved.

4

Small businesses feel the impact of big data too. According to research presented by the Aberdeen Group:

45% feel data formats are inaccessible or underused in analysis.

39% feel data is too "siloed" and information is not shared across applications and business areas.

51% feel it takes too long to get information.

And 35% feel data volume is growing too rapidly.

The "Big" aspect of Big Data is often in the eye of the beholder and for small companies terabyte or even gigabyte scale data can pose the same problems and opportunities that petabytes do for global enterprises.

To compete in today's environment small businesses need more than your everyday business management solution.

What SAP HANA brings to SMEs



Agility: Get real-time business information at the moment you need it so you can clearly define and focus on the right priorities.



Insight: Leverage a single platform for analytics and transactions to get unprecedented insight-to-action capabilities. Leverage pre-delivered apps to solve “un-solvable” problems.



Efficiency: Empower employees with information search capabilities and interactive analysis tools to help them become more efficient and independent of IT staff.



Value: Maintain a simplified IT landscape with a solution priced for small businesses and designed for scalability.

Here's what SAP HANA brings to small and medium enterprises:

Agility: You can get real-time business information at the moment you need it so you can clearly define and focus on the right priorities.

Insight: Leverage a single platform for analytics and transactions to get unprecedented insight-to-action capabilities. Leverage pre-delivered apps to solve “unsolvable” problems.

Efficiency: Empower employees with information search capabilities and interactive analysis tools to help them become more efficient and independent of IT staff.

Value: Maintain a simplified IT landscape with a solution priced for small businesses and designed for scalability.

SAP HANA underlies SAP Strategy



Extending
leadership in
Applications



Broadening
Analytics
footprint



Strengthening
leadership in
Mobile



Becoming
fastest growing
D&T company



Becoming a
profitable **Cloud**
leader

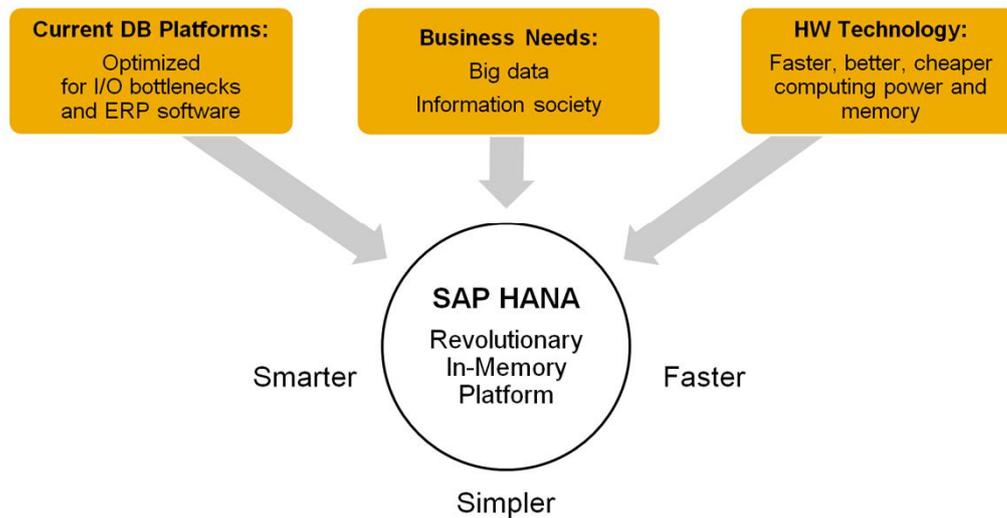
Innovation powered by SAP HANA

Deliver the highest customer value through innovation without disruption

SAP's strategy focuses on 5 main areas: applications, analytics, mobile, database and technology and the cloud. SAP HANA underlies all five of these areas. It underlies our total innovation drive. It is the root for everything that SAP as a company will provide for innovation.

As a partner you will benefit from the innovations available from HANA and it will touch all the parts of your business from sales, through implementation and your development and IT infrastructure. With SAP HANA as the base you will be able to deliver the highest customer value through innovation without disruption.

HANA – In-Memory Platform



© 2014 SAP AG. All rights reserved.

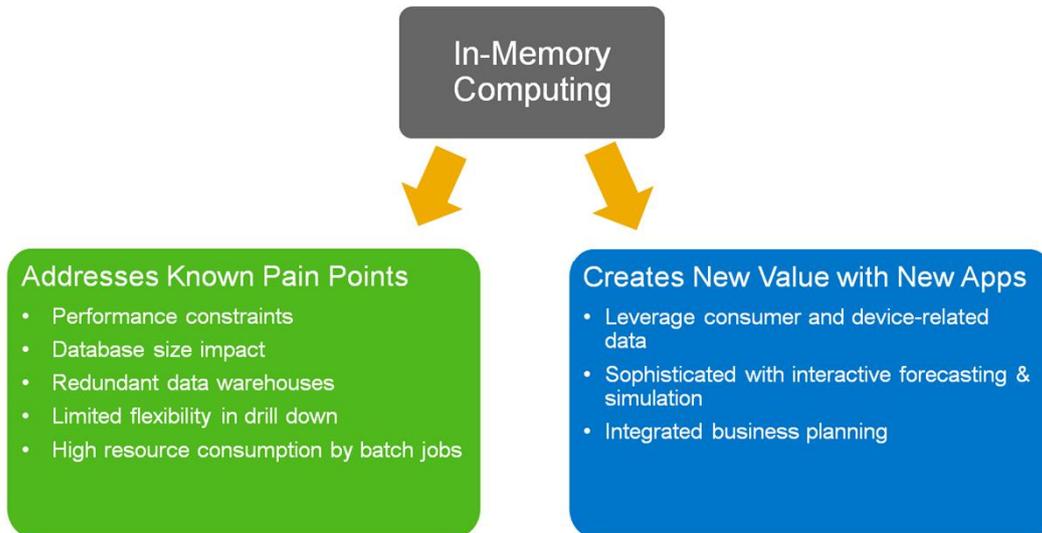
7

SAP HANA is a revolutionary in memory platform brought to you by SAP. It brings together the advances in current database platforms and hardware technology to meet today's business needs.

We can leverage the power of SAP HANA in-memory computing to help small businesses become smarter and faster and run simpler in order to develop competitive advantage.

At its core HANA uses an innovative in-memory technique to store data in a way that is particularly suited for efficiently handling growing amounts of tabular, or relational, data with unprecedented performance.

In-Memory Technology drives Business Opportunities



Technology trends have dramatically lowered the cost of main memory. In memory computing addresses the known pain points for businesses: performance constraints, database size, the need to keep redundant data in data warehouses, limited flexibility to drill down in reports, and high resource consumption by batch jobs. In-memory computing creates new value by allowing us to create new applications that leverage consumer and device-related data, apps that are sophisticated with interactive forecasting and simulation, and integrated business processing. The ability to ask any question (without knowing it in advance) on any data (without knowing the data format in advance) is now the edge the companies need from their technology infrastructure.

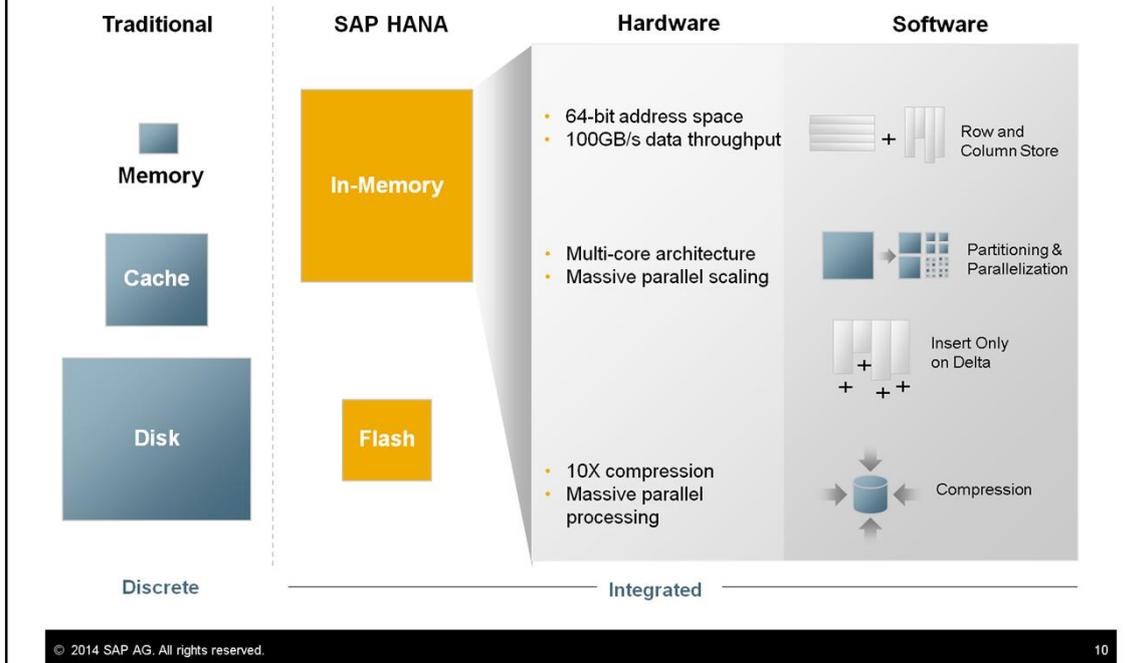
In-Memory Computing

- ❑ SAP HANA caches data and calculates operations in memory.
- ❑ Hard disks are only needed to record changes for permanent persistency.
- ❑ Queries can be executed rapidly and in parallel
- ❑ Coding techniques used in data warehouses are no longer needed.



As SAP HANA caches data in memory and does all calculations in memory, hard disks are only needed to record changes to the database for permanent persistency. The SAP HANA database conceptually is about increasing speed, increasing execution speed of database queries via the use of in-memory data storage, and also increasing speed of application development. With the SAP HANA database, queries can be executed rapidly and in parallel. This means that complex coding techniques, for example, pre-calculation of values that were required by traditional databases to maintain performance are no-longer needed. This is because you can use SAP HANA to query even the biggest datasets in real-time.

In-Memory Computing - Overview



Here we contrast a traditional database with SAP HANA. SAP HANA was designed to run on modern distributed computers built out of multi-core CPUs. It takes advantage of the latest advances in hardware and software.

SAP HANA keeps all data available in main memory. Required backups take place asynchronously as a background task to either disk or solid state drives, so no time is lost to reading data from and saving data to a disk.

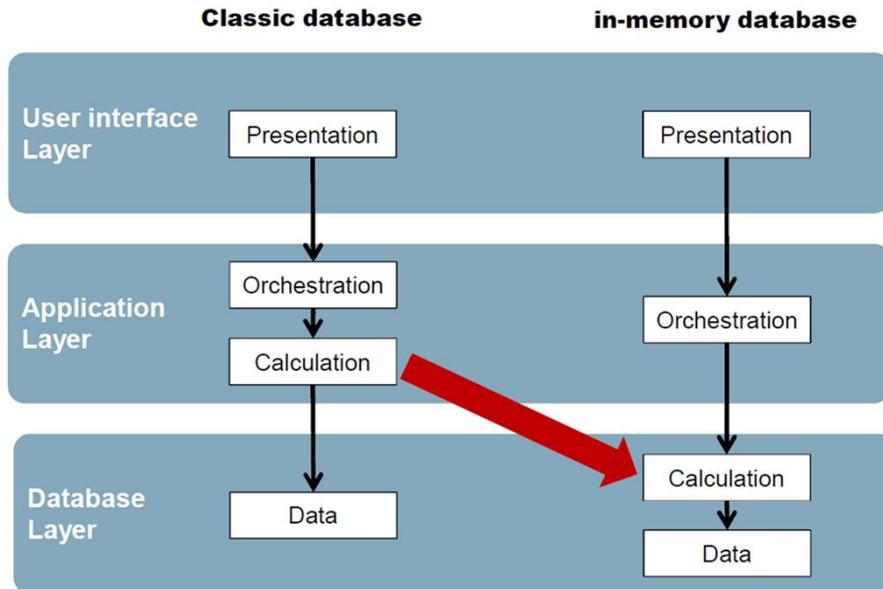
Using columnar data in addition to storing data in rows optimizes searching and aggregations.

SAP HANA was designed from the ground up to perform its basic calculations (such as analytic joins, scans and aggregations) in parallel, fully utilizing the available computing resources of distributed systems.

Additionally, SAP HANA keeps the number of changes to your dataset as small as possible by recording every change as delta to your original dataset. As all of the old data is retained, your applications can effectively provide views of the data as it has changed over time.

SAP HANA takes advantage of compression as well. Compressed data can be loaded into the CPU cache faster. This is because the limiting factor is the data transport between memory and CPU cache, and so the performance gain will exceed the additional computing time needed for decompression.

Minimize Data Movement



© 2014 SAP AG. All rights reserved.

11

Another advantage of using in-memory technology is that it minimizes data movement.

A classic database performs calculations in the application layer. SAP HANA performs calculations in the database layer.

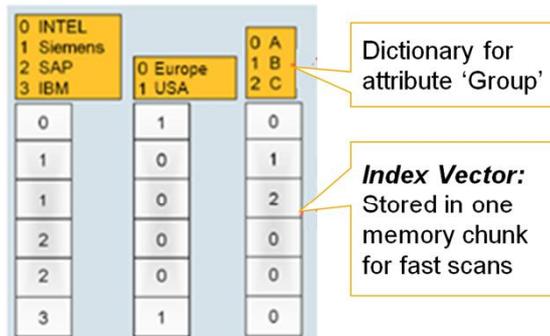
Column Store with SAP HANA

Traditional Database

Company [CHAR50]	Region [CHAR30]	Group [CHAR5]
INTEL	USA	A
Siemens	Europe	B
Siemens	Europe	C
SAP	Europe	A
SAP	Europe	A
IBM	USA	A

- Traditional databases store records in rows
- Columnar layout supports sequential memory access

SAP HANA Database



- Distinct field values are encoded and table rows are only index vectors
- **Result:** faster record scanning (search), better compression, fast aggregation, more efficient join operations

One of the cornerstones of SAP HANA is the way it stores the data in columns. While the concept of columnar data organization is not that new, the fact that SAP HANA delivers a full blown database management system based on that concept is a novelty. But let us take a moment to familiarize ourselves with the concept of column store.

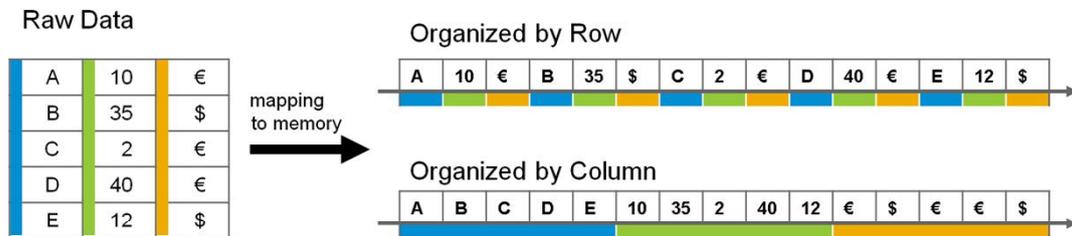
We all probably know that traditional databases store the records in rows. So whenever we need to access certain information, we ask the database to give back certain rows of data which fulfill our query. But this approach is not always the most efficient one. In fact, in many cases we only require certain fields of the original record, or we do not even care about the actual the records themselves, we just would like to know for example, what were the total sales for a certain product group in the last month. In this case it is very inefficient that the database has to fetch all rows for last month, also processing fields that are not relevant to the original query, such as material number and other attributes.

In this particular, and in many other cases, it is more sensible to store the fields and the distinct field values instead of the rows themselves.

SAP HANA creates a dictionary of each distinct field by column, and the table rows become only so-called index vectors.

The results are much faster search, better compression, fast aggregation, and more efficient join operations.

Data Optimization



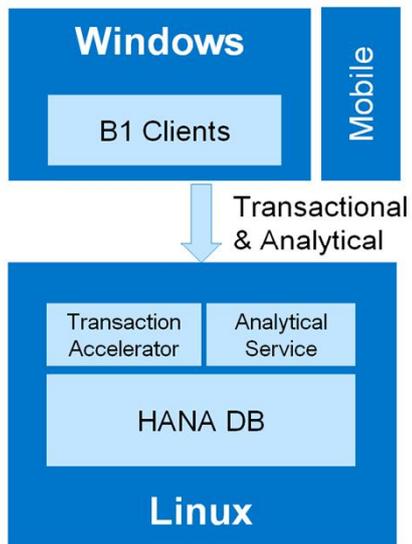
- SAP HANA supports data storage in columns and rows, and is optimized for column order storage.
- Columnar data storage allows for highly efficient compression.
- SAP HANA allows you to read around unwanted data by this organization.

On the graphic, you can get a glimpse about the difference in how data is organized in the traditional row store and the way SAP HANA stores data in a columnar fashion.

SAP HANA supports both data storage in columns and rows, and is optimized for column-order storage. Columnar data storage allows for highly efficient compression.

SAP HANA allows you to read around unwanted data by this organization. This means when you ask for the quantity of sales by country, you do not have to wait for the database to fetch data that it does not need as all the data in a table column is stored in an adjacent manner.

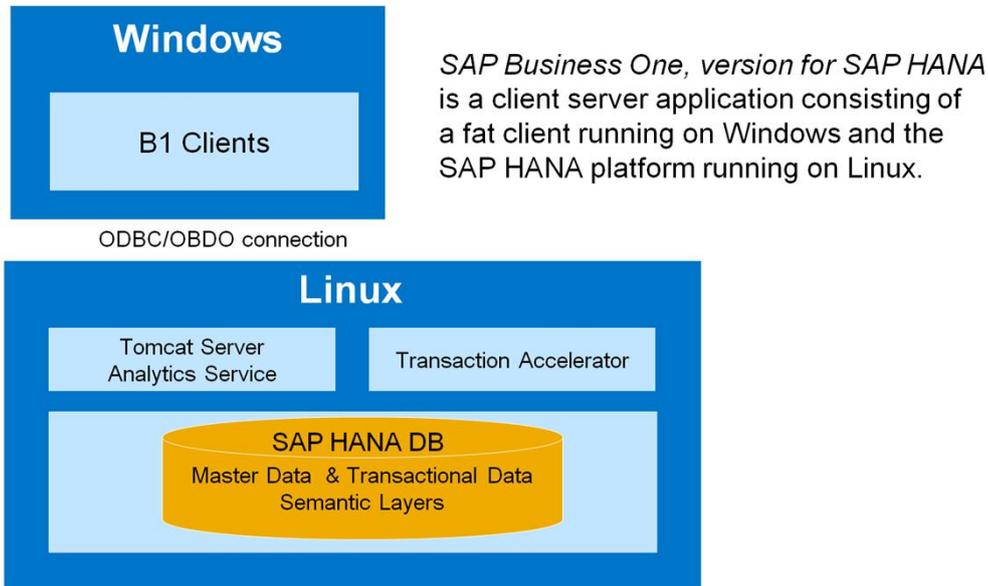
Transactions & Analytics Together



- Analytics was traditionally separate from transactions
- That spurred growth of data warehouses.
- Now we can bring analytics and transactions back together.
- SAP Business One, version for SAP HANA: One database for **both** transactions and analytics.

In the past, analytics was traditionally separate from transactions because of the speed of the disk. We had to set up analytics to answer predefined questions. That spurred growth of data warehouses. Now with in-memory technology we can bring analytics and transactions back together. With SAP Business One, version for SAP HANA, there is one database for **both** online transaction processing and online analytical processing.

Architecture



© 2014 SAP AG. All rights reserved.

15

In terms of architecture, SAP Business One, version for SAP HANA is a client-server application consisting of a client running on Windows and the SAP HANA platform running on a Linux box.

The entire SAP Business One database runs on the SAP HANA database server. The SAP HANA database not only stores data, but also uses triggers as well as views, especially for reporting and upgrade purposes. The SAP HANA database server runs on the Linux box. The SAP HANA database contains the master data and transaction data for all SAP Business One company databases and the SBO common database. The SAP HANA database also contains the semantic layers. Semantic layers are views that are structured for reporting.

Also in the Linux box, we have a Tomcat Server with the analytics service, and the transaction accelerator. The analytics service has the web tier for presentation and business tier services. The transaction accelerator is used for transaction optimization and contains the extreme applications that use SAP HANA's power to extend areas such as availability to promise and cash flow forecasting.

The SAP Business One client applications run in the windows environment. The connection to the Linux server is by ODBC/OBDO connection. The Tomcat server connects to the web browsers for analytics and search.

SAP Business One, version for SAP HANA

- ❑ Fast Google-like freestyle **enterprise search**
- ❑ Pre-delivered **dashboards and reports** optimized for SAP HANA
- ❑ Powerful **interactive analysis reporting** tool with pre-delivered semantic layers
- ❑ **SAP HANA engine and studio** allowing your to customize and/or build your own semantic layers
- ❑ **Innovative out-of-the-box extreme applications, such as:** Cash flow forecasting, Advanced ATP and the Dashboard Designer



© 2014 SAP AG. All rights reserved.

16

Because SAP Business One is running on SAP HANA, it contains features not possible without its groundbreaking technology.

- The fast Google-like freestyle **enterprise search**
- Pre-delivered **dashboards and reports** optimized for SAP HANA
- Powerful **interactive analysis reporting** tool with pre-delivered semantic layers
- **SAP HANA engine and studio** allowing your to customize and/or build your own semantic layers, and
- **Innovative out-of-the-box extreme applications**, such as Cash-flow forecasting, advanced ATP, and the dashboard designer for pervasive analysis.

These extreme apps are just the beginning. SAP HANA provides the foundation for more innovative applications that take advantage of an in-memory database and calculation engine, allowing customers to conduct complex planning, forecasting, and simulation based on real-time data.

In the next unit we will look at each of these features.

Summary



Key points:

- SAP Business One consists of one SAP HANA database for both online transaction processing and for online analytical processing.
- SAP HANA uses an innovative in-memory technique for handling very large amounts of data with unprecedented performance.
- SAP HANA supports both data storage in columns and rows, but is optimized for column-order storage.
- SAP Business One clients run on Windows and the SAP HANA database runs on Linux.

Here are some key points to take away from this course.

- SAP Business One consists of one SAP HANA database for both online transaction processing and for online analytical processing.
- SAP HANA uses an innovative in-memory technique for handling very large amounts of data with unprecedented performance.
- SAP HANA supports both data storage in columns and rows, but is optimized for column-order storage.
- SAP Business One clients run on Windows and the SAP HANA database runs on Linux.
- The Linux box contains the HANA database and a Tomcat server.

Thank you

This concludes the Overview of SAP
Business One version for SAP HANA.

Thank you for your time.

This concludes the overview of SAP Business One version for SAP HANA. Thank you for your time.

© 2014 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Excel, Outlook, PowerPoint, Silverlight, and Visual Studio are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, z10, z/VM, z/OS, OS/390, zEnterprise, PowerVM, Power Architecture, Power Systems, POWER7, POWER6+, POWER6, POWER, PowerHA, pureScale, PowerPC, BladeCenter, System Storage, Storwize, XIV, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, AI/X, Intelligent Miner, WebSphere, Tivoli, Informix, and Smarter Planet are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the United States and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are trademarks or registered trademarks of Adobe Systems Incorporated in the United States and other countries.

Oracle and Java are registered trademarks of Oracle and its affiliates.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems Inc.

HTML, XML, XHTML, and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Apple, App Store, iBooks, iPad, iPhone, iPhoto, iPod, iTunes, Multi-Touch, Objective-C, Retina, Safari, Siri, and Xcode are trademarks or registered trademarks of Apple Inc.

IOS is a registered trademark of Cisco Systems Inc.

RIM, BlackBerry, BBM, BlackBerry Curve, BlackBerry Bold, BlackBerry Pearl, BlackBerry Torch, BlackBerry Storm, BlackBerry Storm2, BlackBerry PlayBook, and BlackBerry App World are trademarks or registered trademarks of Research In Motion Limited.

Google App Engine, Google Apps, Google Checkout, Google Data API, Google Maps, Google Mobile Ads, Google Mobile Updater, Google Mobile, Google Store, Google Sync, Google Updater, Google Voice, Google Mail, Gmail, YouTube, Dalvik and Android are trademarks or registered trademarks of Google Inc.

INTERMEC is a registered trademark of Intermec Technologies Corporation.

Wi-Fi is a registered trademark of Wi-Fi Alliance.

Bluetooth is a registered trademark of Bluetooth SIG Inc.

Motorola is a registered trademark of Motorola Trademark Holdings LLC.

Computop is a registered trademark of Computop Wirtschaftsinformatik GmbH.

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP BusinessObjects Explorer, StreamWork, SAP HANA, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects Software Ltd. Business Objects is an SAP company.

Sybase and Adaptive Server, iAnywhere, Sybase 365, SQL Anywhere, and other Sybase products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Sybase Inc. Sybase is an SAP company.

Crossgate, m@gic EDDY, B2B 360°, and B2B 360° Services are registered trademarks of Crossgate AG in Germany and other countries. Crossgate is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

The information in this document is proprietary to SAP. No part of this document may be reproduced, copied, or transmitted in any form or for any purpose without the express prior written permission of SAP AG.